

AMENDMENTS TO THE SPECIFICATION

Kindly amend the specification as follows:

Kindly replace the paragraph on page 1 beginning with the words “Routing of packets” with the following paragraph:

Routing of packets in a connectionless computer network is now described by way of example with reference to Fig. 1. When a node A (100) sends a packet to a node B (102), A (100) must specify the address of B (102) as the destination address of the packet. The first router R1 (108) that accepts the packet forwards the packet to the next router R2 (110) on the path to B (102), whereupon R2 (110) forwards the packet to the next router R3 (112) on the path to B (102), and so on. When the packet reaches the router to which B (102) is directly connected, it is forwarded to B (102). It may thus be seen that, for any given destination address to which a packet is addressed, every router in the network should know the packet’s next “hop,” i.e., to which next router the packet is to be forwarded. Each router typically maintains this information in a routing table, shown as routing tables 116 and 118, which contains a mapping between addresses or address groups, such as IP subnets, and the next hop for packets destined for these addresses.

Kindly replace the paragraph on page 1 beginning with the words “When a link” with the following paragraph:

When a link connecting two routers in a network fails, a partitioning of the network may occur. Thus in Fig. 1, if the link between R1 (108) and R2 (110) fails, nodes A (100) and C (104) can still communicate with each other but not with nodes B (102) and D (106), and vice versa. Each router will typically automatically detect this situation and update its routing table accordingly, such as by eliminating entries whose next hop is unreachable. However, nodes in one partition may still try to send packets to nodes in the other partition. When this occurs, a “no route to destination” error is typically generated and logged by the first router to detect the problem, which then reports the problem to the

network management system (NMS). The NMS must then decide what action to take, such as tracing the error to its root cause. In large networks where there may be many active communication sessions between nodes at one time, a single link failure event might cause numerous “no route to destination” notifications to be generated in every router in one partition which receives packets that are destined for the other partition and reported to the NMS. Thus, where the existence of a link failure is already known to the NMS, it would be advantageous to know whether or not a routing error is caused by the link failure, as well as which nodes might be affected by the link failure, obviating the need for the NMS to take action that it would normally take.

Kindly replace the paragraph on page 6 beginning with the words “A network management system” with the following paragraph:

A network management system (NMS) 210 preferably maintains copies of routing tables 202 and 204. Having detected a link failure between R1 and R2 (Fig. 3, step 300), NMS 210 may create a connectivity table 212 indicating which nodes are in each of partitions 206 and 208 (Fig. 3, step 302). Since NMS 210 knows that R2 is inaccessible to R1 via link 200, NMS 210 may associate with partition 206 those node addresses in its copy of routing table 202 whose next hop is R2. Likewise, NMS 210 may associate with partition 208 those node addresses in routing table 204 whose next hop is R1. Should NMS 210 receive a “no route to destination” error notification (Fig. 3, step 304) from a network router together with the source and destination addresses of the packet that could not be delivered, NMS 210 may look up the source and destination addresses in connectivity table 212 (Fig. 3, step 306) to determine whether they are from different partitions. If both the source and destination addresses are from different partitions (Fig. 3, step 308), then the “no route to destination” error notification may be an attempt to send the packet across failed link 200. Thus, the error notification may be correlated with the link failure (Fig. 3, step 310) that is already known to NMS 210, and the error may be suppressed and need not be investigated further. Alternatively (Fig. 3, step 312), the error notification should not be

correlated with the link failure and may be investigated or otherwise acted upon by NMS 210.

Kindly replace the paragraph on page 7 beginning with the words “Reference is now made to Fig. 4” with the following paragraph:

Reference is now made to Fig. 4, which is a simplified flowchart illustration of a method of correlation of routing errors to link failures in a connectionless network supported by a distributed network management system, operative in accordance with a preferred embodiment of the present invention. In Fig. 4 the present invention is implemented in a distributed network management system, such as is described in U.S. Patent Application No. 09/799,637 and published as Published Application No. 20010039577, where every router has an associated software agent which continuously monitors the state of the router and its links. The agents monitoring R1 and R2 would thus detect the failure of link 200 (step 400) and then communicate with each other to create connectivity table 212 (step 402) which may then be provided to the agents of all other routers in the network. Thus, when any router Rx encounters a “no route to destination” error (step 404), its associated agent looks up the source and destination addresses in connectivity table 212 (step 406) to determine whether they are from different partitions (step 408), and action may be taken (step 410) or the error notification ignored (step 412) as described above.

Kindly replace the paragraph on page 7 beginning with the words “Reference is now made to Fig. 5” with the following paragraph:

Reference is now made to Fig. 5, which is a simplified flowchart illustration of a method of identifying nodes that may be affected by link failures in a connectionless network, operative in accordance with a preferred embodiment of the present invention. In Fig. 5 a list of virtual paths in a network is maintained (step 500), where each virtual path represents the traversal of the links, routers, and other network elements comprising the most commonly used and/or most heavily used routes between network nodes, as

determined using any predefined measure of use. The virtual path list may be maintained centrally, such as by NMS 210, or in a distributed manner, such as by one or more agents in a distributed network management system. The virtual path list may be created using any conventional technique, such as by identifying common access patterns in router access lists, analyzing network failure alarms (e.g., packet lost, no route, etc.) to determine traffic flow, and determining network tomography from traffic counter patterns. When a failed link is detected (step 502), each virtual path may be checked using any known technique to determine if it is broken (step 504) and, if so, which nodes and other network elements along the path are affected (step 506). Thereafter, should a “no route to destination” error be encountered (step 508) where the source address of the packet being sent belongs to the node at one end of a virtual path known to have a failed link (step 510), and the packet’s destination address belongs to the node at the other end of the virtual path, the error may be correlated to the failed link and action may be taken (step 512) or suppressed (step 514) as described hereinabove.